

Diskshipping-less operation

Bob Eldering

JIVE

eldering@jive.eu

2015-06-26

- 1 FlexBuff
- 2 Synchronization
- 3 Results

What is a FlexBuff?

- Powerful COTS computer
- Lots of disk drives
- Almost drop-in replacement for Mark5

FlexBuff data organization

- All disks mounted under `/mnt/` as `disk<N>/`
- On all disks a directory is created
- File chunks are created in these directories
- Unix tools to access data

FlexBuff usage

```

flexbuff0: 521G in data files not in VBS format, free/used/total: 19.1T/84.2T/109T
Total 89.7T 276G 6.63T 5.15T 4.65T 1.21T 5.97T 1.66T 28.9T 6.28T 4.98T 3.11T 1.62T 82.6G 4.31T 6.28T 18.2G 1.07T 1.51T
EA0590 12.3T 1.08T 1.09T 1.08T 559G 1.09T 1.08T 930G 1.09T 1.08T 1.06T 1.07T 1.10T
EG087A 14.1T 3.20T 2.25T 1.62T 1.30T 1.20T 1.30T 3.25T
EH027D 4.15T 685G 684G 683G 680G 151G
EH030A 1.99T 1.99T
EL048A 4.97T 276G 480G 434G 508G 582G 531G 300G 534G 490G 528G 428G
EL050C 2.18T 2.18T
EL050D 1.67T 1.67T
EP091A 3.84T 3.84T
EP096A 5.42T 802G 730G 801G 801G 730G 82.6G 801G 801G
EP096B 1.72T 1.72T
ER040 754G 754G
ER043A 2.30T 2.30T
ER043B 2.34T 2.34T
F15L1 131G 131G
FBTEST 18.2G 18.2G
FR022 577G 577G
GA032C 2.97T 2.97T
GA036 4.04T 4.04T
GB075B 4.87T 4.87T
GB075C 4.99T 4.99T
GB076 5.25G 5.25G
GL039 2.60T 2.60T
GP053F 1.18T 1.18T
N1402 1.10T 1.10T
N15C1 708G 708G
N15C2 810G 810G
N15L2 427G 427G
N15M1 37.6G 37.6G
N15Q1 761G 761G
N15X1 847G 436G 411G

aribox: 16.1T in data files not in VBS format, free/used/total: 9.11T/60.6T/73.4T
Total 44.5T 2.65T 1.86T 1.08T 1.34T 1.26T 2.13T 5.70T 1.06T 1.28T 1.67G 183G 914G 2.72T 683G 2.62T 1.57T 3.48T 1.12T 757G 500G 1.02T 1.02T
EG084A 2.10T
EG087A 4.54T 663G 1.30T
ES073 8.20T 638G 1.05T 764G 768G 764G 763G 167G 183G 1.30T
F14C2 201G 201G
FBTEST 18.2G
FT009A 265G
GA035A 6.20T 682G 539G 608G 679G 635G 683G 683G 665G 500G
GA035B 3.70T 284G 754G 325G 824G
GB075B 4.65T 4.65T
GR035 7.87T 1.36T 1.08T 914G 1.37T 1.57T
GR35F 1.02T
N1402 1.05T 1.05T
R1680 4.68T 1.26T 1.28T 1.12T 1.02T
  
```

FUSE file system

```
jops@aribox:/tmp/fuse_mnt$ ls -l es073_*no0001
-rwxrwxrwx 0 root root 17431842272 Oct 16 2014 es073_bd_no0001
-rwxrwxrwx 0 root root 17732323648 Oct 16 2014 es073_ef_no0001
-rwxrwxrwx 0 root root 17432832320 Oct 27 2014 es073_jb_no0001
-rwxrwxrwx 0 root root 17430996792 Oct 17 2014 es073_mc_no0001
-rwxrwxrwx 0 root root 17431846400 Oct 28 2014 es073_nt_no0001
-rwxrwxrwx 0 root root 17399795200 Oct 28 2014 es073_on_no0001
-rwxrwxrwx 0 root root 8597734400 Oct 28 2014 es073_ra_no0001
-rwxrwxrwx 0 root root 17431842272 Oct 19 2014 es073_sv_no0001
-rwxrwxrwx 0 root root 17665132168 Oct 21 2014 es073_tr_no0001
-rwxrwxrwx 0 root root 6931571200 Oct 28 2014 es073_ur_no0001
-rwxrwxrwx 0 root root 17463877824 Oct 22 2014 es073_zc_no0001
```

VBS remove

```
jops@ccsbeta:~$ vbs_remove -h
usage: vbs_remove [-h] -e EXPERIMENT [-f FLEXBUFF [FLEXBUFF ...]]
                [-s STATION [STATION ...]]
```

Remove VBS files from flexbuffs and their references from the database

optional arguments:

```
-h, --help                show this help message and exit
-e EXPERIMENT, --experiment EXPERIMENT
-f FLEXBUFF [FLEXBUFF ...], --flexbuff FLEXBUFF [FLEXBUFF ...]
                          Flexbuff(s) to remove files from, default: all local
                          flexbuffs
-s STATION [STATION ...], --station STATION [STATION ...]
                          Station(s) to remove data for, default: all stations
```

Transfer overview

- Experiment schedule from vlbeer
- Experiment ends, locate file chunks
- Extract recording names and transfer them
- Verify all chunks have been transferred correctly
- Enter the chunks in a database, ready for correlation

Locate file chunks

- `/mnt/disk*/<experiment>_<station>_<scan>/<experiment>_<station>_<scan>.<N>`
- `ssh oper@flexbuff`


```
/mnt/disk26/fbtest_x0_no0001:  
fbtest_x0_no0001.00000008  fbtest_x0_no0001.00000022  
fbtest_x0_no0001.00000015  fbtest_x0_no0001.00000029  
  
/mnt/disk27/fbtest_x0_no0001:  
fbtest_x0_no0001.00000005  
  
/mnt/disk28/fbtest_x0_no0001:  
fbtest_x0_no0001.00000002  fbtest_x0_no0001.00000017  
fbtest_x0_no0001.00000010  fbtest_x0_no0001.00000024  
  
/mnt/disk29/fbtest_x0_no0001:  
fbtest_x0_no0001.00000003  fbtest_x0_no0001.00000018  
fbtest_x0_no0001.00000011  fbtest_x0_no0001.00000025
```

Transfer recordings

- `~oper/.JIVE.allow`
`2014-07-01 00:00:00 2042-07-30 00:00:00 1000`
- In parallel with recording
`runtime = flexbuff_transfer_<Station>`
- Choose destination FlexBuff and port

Network graph

- UDT
- Connection per file chunk
- Locating chunks can take a long time



Verify recordings

- Presence
- Size
- `~oper/.JIVE.transferred`
- `!file_check? 0 : VDIF : 32 : 2015y168d14h17m42.0391s : 1.99197s : 32Mbps : 0 ;`

Error handling

- Retry on error
- Continue where error occurred

Correlate

Configuration

Experiment:

5L2TEST

Sort and filter by:

name (last year)

Profile:

flexbuff te

Lock

Create new

Select data input type:

eVLBI

Forking

File

VSN

FlexBuff (remote)

FlexBuff (local)

Configure

Stations

Ar

Ex

Mc

Ox

Sv

Wb

Bd

Hh

Nt

Sh

T6

Wd

Ef

Jb

O8

Sr

Tr

Zc

Frequency points:

32

Integration time:


2

Cross polarizations:


Using a Mark6 as a FlexBuff

- Different disk mount points:
/mnt/disk/[1234]/[0-7]/
- Mount status

- Parallel recording since session 3, 2014
- Network rate capped at 1000Mbps
- DBBC
- Race condition and memory leak




Fila10G + FlexBuff




Mark5B

- Sign/magnitude swap




Fila10G + FlexBuff




Mark5B

- But:




- Kernel settings
- Monitor packet loss: evlbi?

- Still some packet loss at start



- Memory initialization

- Successful second half



- Correlate against other stations

Questions/comments?

Questions/comments?